arXiv:1407.0645v2 [cs.LO] 18 Dec 2015

Branching Bisimilarity of Normed BPA Processes is in NExpTime

Wojciech Czerwiński[1⋆] and Petr Jančar[2]

[1] University of Warsaw
[2] FEI, Technical University of Ostrava
wczerwin@mimuw.edu.pl, petr.jancar@vsb.cz

Abstract. Branching bisimilarity on normed BPA processes was recently shown to be decidable by Yuxi Fu (ICALP 2013) but his proof has not provided any upper complexity bound. We present a simpler approach based on relative prime decompositions that leads to a nondeterministic exponential-time algorithm; this is close to the known exponential-time lower bound.

Key words: Verification, branching bisimulation equivalence, Basic Process Algebra processes, complexity.

1 Introduction

Similarly as language equivalence in automata theory, bisimulation equivalence (also called bisimilarity) is a fundamental notion in theory of processes. The decidability and complexity questions for bisimilarity on various models of infinite-state systems have been explored in a long list of research papers. (See [14] for an updated overview of a specific area of process rewrite systems.)

One of the basic models is called Basic Process Algebra (BPA), which can be related to context-free grammars in Greibach normal form. The configurations are identified with sequences of variables (nonterminals), and a configuration can change by performing an action (rather than reading a symbol) in which case its leftmost variable is rewritten. The seminal paper by Baeten, Bergstra and Klop [1] showed the decidability of bisimilarity for BPA configurations in the normed case, where each variable can be stepwise rewritten to the empty word; this decidability result was later extended to the whole class BPA [4]. The exploration of complexity culminated by showing that the normed case is, in fact, in PTime [9] (see [5] for the so far best known upper bound). The complexity for the whole class BPA is known to lie between ExpTime [12] and 2-ExpTime (claimed in [2] and explicitly proven in [11]).

In the presence of silent (unobservable) actions the problems become harder. The decidability question for weak bisimilarity of (even normed) BPA configurations is a long-standing open problem; we only know ExpTime-hardness here,

⋆ W. Czerwinski acknowledges a partial support by the Polish National Science Centre grant 2013/09/B/ST6/01575.

already in the normed case [13]. There is a similar long-standing open problem for Basic Parallel Processes, which is the parallel (or commutative) version of BPA. Positive results were recently achieved for a finer version of weak bisimilarity, called *branching bisimilarity*. (It was introduced by van Glabbeek and Weijland [8] as the coarsest equivalence respecting branching time in some sense.) It was shown that branching bisimilarity is decidable on the normed Basic Parallel Processes [6], and Yuxi Fu [7] has shown the decidability for normed BPA configurations; the latter is the starting point of our paper. We also note that the fresh paper [15] shows that these decidability results cannot be essentially extended.

Fu's result for branching bisimilarity on normed BPA is substantially stronger than the previous results dealing with so called totally normed BPA [10, 3]; the proof uses an involved tableau framework (being inspired by [10] and related works) and does not provide any upper complexity bound. Regarding the lower bound, Fu notes that the construction used by Mayr [13] for weak bisimilarity can be easily adapted to yield ExpTime-hardness also for branching bisimilarity on normed BPA.

An important novel ingredient of the decidability proof by Yuxi Fu can be called the *class-change norm* (corresponding to the branching norm in [7]): while the standard norm counts all the steps when a configuration is reduced to the empty one, the class-change norm only counts the steps that change the current equivalence-class. It is not clear how to compute this norm directly but equivalent configurations $\alpha \sim \beta$ must agree on this norm. Another useful fact (also observed by Fu) is that the relation of $\alpha\gamma$ and $\beta\gamma$ (either $\alpha\gamma \sim \beta\gamma$ or $\alpha\gamma \not\sim \beta\gamma$) is determined solely by the redundant variables w.r.t. γ, i.e. by those X for which $X\gamma \sim \gamma$, independently of the string γ itself.

Our contribution is based on introducing the *relative prime decomposition* of configurations; unlike the prime decomposition of variables as, e.g., in the case of bisimilarity of normed BPA [9], we introduce a decomposition related to each possible set R of redundant variables (representing respective suffixes γ). The (relative) equivalence $\alpha \sim_R \beta$ can be then replaced with the equality of prime decompositions $\mathrm{PD}_R(\alpha) = \mathrm{PD}_R(\beta)$. We suggest a nondeterministic exponential-time algorithm that guesses the appropriate decompositions and then verifies their correctness in the sense that the equality of the decompositions w.r.t. the guess is indeed a branching bisimulation. We thus place the branching bisimilarity of normed BPA configurations in NExpTime.

Fu [7] has also shown that the respective "regularity problem" (given a normed BPA configuration, is it branching bisimilar to some unspecified finite-state process?) is decidable. Our approach places this problem in NExpTime as well.

Remark. It seems natural to look for a deterministic exponential-time algorithm that would (deterministically) compute the decompositions, e.g., by proceeding via a certain series of decreasing overapproximations. Nevertheless, this question is left open here.

Structure of the paper. In Section 2 we define the used notions and state the result. Section 3 deals with the class-change norm and related observations that essentially already appeared in [7]. Section 4 then introduces the relative prime decompositions of configurations and proves their uniqueness. Section 5 defines the branching bisimilarity bases and shows how their consistency can be checked. Finally Section 6 contains algorithms for the bisimilarity and regularity problems.

2 Preliminaries, and statements of results

We first recall a general definition of branching bisimilarity, which is then applied to BPA configurations. Throughout the text we add some remarks (in italic) related to weak bisimilarity; they are not needed for understanding the presented result.

By $\mathcal{A}^*$ we denote the set of finite sequences of elements of the set $\mathcal{A}$. By ε we denote the empty sequence, and by $|w|$ the length of $w \in \mathcal{A}^*$. We put $\mathbb{N} = \{0, 1, 2, \dots\}$.

Labelled transition systems. A *labelled transition system*, an *LTS* for short, is a tuple

$$\mathcal{L} = (\mathcal{S}, \mathcal{A}, (\stackrel{a}{\longrightarrow})_{a\in\mathcal{A}})$$

where $\mathcal{S}$ is a set of states (at most countable in our case), $\mathcal{A}$ is a set of *actions* (finite in our case), and $\stackrel{a}{\longrightarrow} \subseteq \mathcal{S} \times \mathcal{S}$ is a set of *transitions labelled with* a. We reserve the symbol

$$\tau \text{ for the (only) } \textit{silent action};$$

the *visible actions* are the elements of $\mathcal{A}_{\text{VIS}} = \mathcal{A} \setminus \{\tau\}$. (If $\tau \in \mathcal{A}$, then $\mathcal{A} = \mathcal{A}_{\text{VIS}} \cup \{\tau\}$, otherwise $\mathcal{A} = \mathcal{A}_{\text{VIS}}$.)

We write $s \stackrel{a}{\longrightarrow} t$ rather than $(s,t) \in \stackrel{a}{\longrightarrow}$ (for $a \in \mathcal{A}$), and we define $s \stackrel{w}{\longrightarrow} t$ for $w \in \mathcal{A}^*$ inductively: $s \stackrel{\varepsilon}{\longrightarrow} s$; if $s \stackrel{a}{\longrightarrow} s'$ and $s' \stackrel{u}{\longrightarrow} t$, then $s \stackrel{au}{\longrightarrow} t$. By $s \stackrel{w}{\longrightarrow} t$ we sometimes also refer to a concrete respective path from s to t in $\mathcal{L}$.

Branching bisimilarity. Given an LTS $\mathcal{L} = (\mathcal{S}, \mathcal{A}, (\stackrel{a}{\longrightarrow})_{a\in\mathcal{A}})$, a *symmetric* relation $\mathcal{B} \subseteq \mathcal{S} \times \mathcal{S}$ is a *branching bisimulation* if for any $(s,t) \in \mathcal{B}$, $a \in \mathcal{A}$, and $s' \in \mathcal{S}$ such that $s \stackrel{a}{\longrightarrow} s'$ we have:

- $a = \tau$ and $(s', t) \in \mathcal{B}$, or
- ($a \in \mathcal{A}_{\text{VIS}} \cup \{\tau\}$ and) there is a sequence $t = t_0 \stackrel{\tau}{\longrightarrow} t_1 \stackrel{\tau}{\longrightarrow} \cdots \stackrel{\tau}{\longrightarrow} t_k \stackrel{a}{\longrightarrow} t'$ (for some $k \geq 0$) such that $(s', t') \in \mathcal{B}$ and $(s, t_i) \in \mathcal{B}$ for all $i \in \{1, 2, \dots, k\}$.

By $s \sim t$, to be read as "states s, t are *branching bisimilar*", we denote that there is a branching bisimulation containing (s,t). We can easily verify the standard facts that $\sim$ is the union of all branching bisimulations, and thus the maximal branching bisimulation, and that $\sim$ is an equivalence relation.

Remark 1. Weak bisimulations, *and* weak bisimilarity, *are defined as above but we allow sequences with "post" τ-transitions, like $t = t_0 \xrightarrow{\tau} t_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} t_k \xrightarrow{a} t'_0 \xrightarrow{\tau} t'_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} t'_\ell = t'$ (for $\ell \geq 0$), and we only require that the final pair (s', t') belongs to $\mathcal{B}$. We denote the weak bisimilarity by $\approx$. Any branching bisimulation is thus a weak bisimulation; hence branching bisimilarity $\sim$ is finer than weak bisimilarity $\approx$. They coincide in the case with no silent action, in which case we use the notion of (strong) bisimilarity. In the system given by the following transitions we have $s_1 \not\sim s_2$ but $s_1 \approx s_2$: $s_1 \xrightarrow{\tau} s_2$, $s_1 \xrightarrow{a} s_5$, $s_2 \xrightarrow{\tau} s_3$, $s_3 \xrightarrow{a} s_5$, $s_2 \xrightarrow{a} s_4$, $s_4 \xrightarrow{b} s_5$.*

Normed BPA systems. A *BPA system* is given by a context-free grammar in Greibach normal form, with no starting variable (nonterminal). We denote it as

$$\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$$

where $\mathcal{V}$ is a finite set of *variables* (or nonterminals), $\mathcal{A}$ is a finite set of *actions* (or terminals), which can contain the *silent action* τ, and $\mathcal{R}$ is a finite set of *rules* of the form $A \xrightarrow{a} \alpha$ where $A \in \mathcal{V}$, $a \in \mathcal{A}$, $\alpha \in \mathcal{V}^*$.

A BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ has the associated LTS

$$\mathcal{L}_\mathcal{G} = (\mathcal{V}^*, \mathcal{A}, (\xrightarrow{a})_{a \in \mathcal{A}})$$

where each rule $A \xrightarrow{a} \alpha$ in $\mathcal{R}$ induces the transitions $A\beta \xrightarrow{a} \alpha\beta$ for all $\beta \in \mathcal{V}^*$. The states of $\mathcal{L}_\mathcal{G}$, i.e. the strings of variables, are called *configurations* or *processes*.

A *variable* $A \in \mathcal{V}$ is *normed* if there is $w \in \mathcal{A}^*$ such that $A \xrightarrow{w} \varepsilon$. A *BPA system* $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ is *normed*, i.e. an *nBPA system*, if each $A \in \mathcal{V}$ is normed.

Branching bisimilarity problem for nBPA. By the *branching bisimilarity problem for normed BPA* we mean the decision problem specified as follows:

> *Instance:* a normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ and variables $A, B \in \mathcal{V}$.
> *Question:* Is $A \sim B$ in $\mathcal{L}_\mathcal{G}$?

The variant with general configurations $\alpha, \beta \in \mathcal{V}^*$ in the instances, asking whether $\alpha \sim \beta$, can be easily reduced to the above variant with variables A, B.

Semantic finitess (or regularity) problem. In our context, the *regularity problem* is specified as follows:

> *Instance:* a normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ and $\alpha \in \mathcal{V}^*$.
> *Question:* Is $\alpha \sim s$ for a state s in some finite LTS ?

As usual, when comparing states in two different LTSs $\mathcal{L}_1, \mathcal{L}_2$, we implicitly refer to the LTS arising as the disjoint union of $\mathcal{L}_1$ and $\mathcal{L}_2$.

Results. The next two theorems capture our main results.

Theorem 2. *The branching bisimilarity problem for normed BPA is in* NExpTime.

Theorem 3. *The regularity problem for normed BPA (w.r.t. branching bisimilarity) is in* NExpTime.

We prove the theorems in the following sections; in the rest of this section we recall some facts about the standard norm, and we provide an example.

Standard norm. Given a normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, the norm $\|\alpha\|$ of $\alpha \in \mathcal{V}^*$ is the length $|w|$ of a shortest $w \in \mathcal{A}^*$ such that $\alpha \xrightarrow{w} \varepsilon$. (Note that the silent steps $\xrightarrow{\tau}$ are also counted.) A *transition* $\alpha \xrightarrow{a} \beta$ is *norm-reducing* if $\|\alpha\| > \|\beta\|$, in which case $\|\beta\| = \|\alpha\| - 1$, in fact.

The facts captured by the next proposition are standard and easy; they also entail that we can check in polynomial time whether a BPA system is normed.

Proposition 4.
(1) $\|\varepsilon\| = 0$.
(2) $\|\alpha\beta\| = \|\alpha\| + \|\beta\|$.
(3) $\|A\| = 1 + \|\alpha\|$ *for a (norm-reducing) rule* $A \xrightarrow{a} \alpha$.
(4) There is a polynomial-time algorithm (based on dynamic programming) that computes $\|A\|$ *for each* $A \in \mathcal{V}$ *(when given* $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$*).*
(5) The values $\|A\|$ *are exponentially bounded (in the size of* $\mathcal{G}$*).*

Example 5. *Let* $\mathcal{V} = \{S_1, S_2, S_3\} \cup \{M_R \mid \emptyset \neq R \subseteq \{1,2,3\}\} \cup \{A, B, C, D\}$ *and* $\mathcal{A} = \{a_1, a_2, a_3\} \cup \{\tau\}$. *We write just* M_{13} *instead of* $M_{\{1,3\}}$, *etc. Let* $\mathcal{R}$ *be the set containing the following rules:*

$$S_1 \xrightarrow{a_1} \varepsilon,\ S_1 \xrightarrow{\tau} \varepsilon,\ S_2 \xrightarrow{a_2} \varepsilon,\ S_2 \xrightarrow{\tau} \varepsilon,\ S_3 \xrightarrow{a_3} \varepsilon,\ S_3 \xrightarrow{\tau} \varepsilon,$$

$$M_R \xrightarrow{a_i} M_R,\ M_R \xrightarrow{\tau} \varepsilon \text{ for all nonempty } R \subseteq \{1,2,3\} \text{ and } i \in R,$$

$$A \xrightarrow{\tau} S_1 M_3,\ B \xrightarrow{a_1} C,\ B \xrightarrow{\tau} M_2 M_3,\ C \xrightarrow{a_1} C,\ C \xrightarrow{\tau} M_3 M_2.$$

Here $\|S_i\| = 1, \|M_R\| = 1, \|A\| = \|B\| = \|C\| = 3$. *We can check that* $M_{23} \not\sim M_3 M_2$, *and* $S_2 M_{23} \sim M_{23} \sim M_3 M_{23}$, *though* $\|M_{23}\| \neq \|M_3 M_{23}\|$.

3 Class-change norm, and relative equivalences

In this section we recall some standard facts, the special norm introduced by Y. Fu (given by a slightly modified definition here), and some observations presented already in [7]. The main notions introduced here are the *class-change norm* and the *relative equivalences* $\sim_R$ for subsets R of the set of variables. We implicitly refer to a given normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, though some claims hold more generally. We give all proofs, to be self-contained.

Congruence property, and silent variables.

Proposition 6. *If $\alpha \sim \beta$ and $\gamma \sim \delta$, then $\alpha\gamma \sim \beta\delta$.*

Proof. It is trivial to check that $\{(\alpha\gamma, \alpha\delta) \mid \alpha, \gamma, \delta \in \mathcal{V}^*, \gamma \sim \delta\}$ is a branching bisimulation. Hence $\gamma \sim \delta$ implies $\alpha\gamma \sim \alpha\delta$. Slightly more subtle is to check that $\{(\alpha\gamma, \beta\gamma) \mid \alpha, \beta, \gamma \in \mathcal{V}^*, \alpha \sim \beta\}$ is a branching bisimulation. □

The second point in the above proof depends on our normedness assumption. (E.g., for the unnormed variable A with the only rule $A \xrightarrow{\tau} A$ we have $\varepsilon \sim A$ but $\delta \not\sim A\delta$ if $\delta \not\sim \varepsilon$, since $A\delta \sim \varepsilon$.)

We say that $A \in \mathcal{V}$ is a *silent variable* if $A \xrightarrow{w} \alpha$ implies $w \in \{\tau\}^*$ (i.e., we can never perform a visible action when starting from A); let $\mathcal{V}_{\text{SIL}}$ consist of all silent variables. We observe that $\varepsilon \sim \alpha$ iff $\alpha \in (\mathcal{V}_{\text{SIL}})^*$ (in our normed case); hence $\alpha \in (\mathcal{V}_{\text{SIL}})^*$ implies $\gamma \sim \alpha\gamma$ for any $\gamma \in \mathcal{V}^*$.

Convention (*on silent variables*). Since the silent variables can be determined by a straightforward polynomial algorithm and they can be removed from any $\beta \in \mathcal{V}^*$ without changing its equivalence class, we further assume that our normed BPA systems have no silent variables.

Class-change norm. Example 5 showed that we can have $\alpha \sim \beta$ though $\|\alpha\| \neq \|\beta\|$. We now define a norm for which this cannot happen.

A transition $\alpha \xrightarrow{a} \beta$ is *class-changing* if $\alpha \not\sim \beta$. The *cc-length* of a path $\alpha \xrightarrow{w} \beta$ is the number of class-changing transitions in the path. We note that the cc-length of $\alpha \xrightarrow{u} \beta \xrightarrow{v} \gamma$ is the sum of the cc-lengths of $\alpha \xrightarrow{u} \beta$ and $\beta \xrightarrow{v} \gamma$.

The *class-change norm* (or the *cc-norm*) of $\alpha \in \mathcal{V}^*$, denoted $\langle\!\langle \alpha \rangle\!\rangle$,

is the minimum of the cc-lengths of paths $\alpha \xrightarrow{w} \varepsilon$.

Any shortest path $\alpha \xrightarrow{w} \varepsilon$ such that its cc-length is equal to $\langle\!\langle \alpha \rangle\!\rangle$ is called

a *witness path for* α.

We observe that if $\alpha \xrightarrow{u} \beta \xrightarrow{v} \varepsilon$ is a witness path for α then $\beta \xrightarrow{v} \varepsilon$ is a witness path for β.

Remark 7. Consider the rules $A \xrightarrow{\tau} A'$, $A' \xrightarrow{\tau} \varepsilon$, $A \xrightarrow{a} \varepsilon$, $B \xrightarrow{a} B$, $B \xrightarrow{b} \varepsilon$. Here $AB \sim A'B \sim B$, and the only witness path for AB is $AB \xrightarrow{a} B \xrightarrow{b} \varepsilon$, with the cc-length 1; hence $\langle\!\langle AB \rangle\!\rangle = 1$ (while $\|AB\| = 2$). The branching norm in [7] would be also 1, but witnessed by $AB \xrightarrow{\tau} A'B \xrightarrow{\tau} B \xrightarrow{b} \varepsilon$.

A transition $\alpha \xrightarrow{a} \beta$ is *class-change-norm-reducing*, *ccn-reducing* for short, if $\langle\!\langle \alpha \rangle\!\rangle > \langle\!\langle \beta \rangle\!\rangle$, i.e., if $\langle\!\langle \beta \rangle\!\rangle = \langle\!\langle \alpha \rangle\!\rangle - 1$.

Proposition 8.
(1) $\langle\!\langle \alpha \rangle\!\rangle \leq \|\alpha\|$.
(2) If $\alpha \sim \beta$, then $\langle\!\langle \alpha \rangle\!\rangle = \langle\!\langle \beta \rangle\!\rangle$.

Proof. (1) is trivial.
(2): Informally speaking, any class-changing transition must be matched by a response that finishes by a corresponding class-change, and thus the claim is intuitively clear. Formally we suppose a counterexample

$$\alpha \sim \beta,\ \langle\!\langle \alpha \rangle\!\rangle < \langle\!\langle \beta \rangle\!\rangle,$$

where α has the shortest possible witness path $\alpha \xrightarrow{w} \varepsilon$. We cannot have $\alpha = \varepsilon$, since in this case $\alpha = \beta = \varepsilon$ (recall that we have excluded silent variables), and $\langle\!\langle \alpha \rangle\!\rangle = \langle\!\langle \beta \rangle\!\rangle = 0$. Hence $\alpha = A\delta$, and the path $\alpha \xrightarrow{w} \varepsilon$ can be written as

$$\alpha = A\delta \xrightarrow{a} \gamma\delta = \alpha' \xrightarrow{w'} \varepsilon$$

(for a rule $A \xrightarrow{a} \gamma$). Obviously, $\alpha' = \gamma\delta$ has a shorter witness path than α, and $\langle\!\langle \alpha' \rangle\!\rangle \leq \langle\!\langle \alpha \rangle\!\rangle < \langle\!\langle \beta \rangle\!\rangle$. Our assumptions thus imply $\alpha' \not\sim \alpha$ (otherwise $\alpha' \sim \beta$ with $\langle\!\langle \alpha' \rangle\!\rangle < \langle\!\langle \beta \rangle\!\rangle$ would constitute a "smaller" counterexample); hence $\langle\!\langle \alpha' \rangle\!\rangle = \langle\!\langle \alpha \rangle\!\rangle - 1$. Since $\alpha \sim \beta$, there is a response to the transition $\alpha \xrightarrow{a} \alpha'$ (i.e., to $A\delta \xrightarrow{a} \gamma\delta$), namely a sequence

$$\beta = \beta_0 \xrightarrow{\tau} \beta_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k \xrightarrow{a} \beta'$$

(for some $k \geq 0$), such that $\alpha' \sim \beta'$ and $\alpha \sim \beta_i$ for all $i \in \{0, 1, \ldots, k\}$; thus the transitions in $\beta_0 \xrightarrow{\tau} \beta_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k$ are not class-changing. Hence $\langle\!\langle \beta \rangle\!\rangle \leq \langle\!\langle \beta' \rangle\!\rangle + 1$, and thus $\langle\!\langle \beta' \rangle\!\rangle \geq \langle\!\langle \beta \rangle\!\rangle - 1 > \langle\!\langle \alpha \rangle\!\rangle - 1 = \langle\!\langle \alpha' \rangle\!\rangle$; this implies that α', β' (where $\alpha' \sim \beta'$ and $\langle\!\langle \alpha' \rangle\!\rangle < \langle\!\langle \beta' \rangle\!\rangle$) constitute a smaller counterexample. □

In Example 5 we can check that $\langle\!\langle S_1 M_{12} \rangle\!\rangle = 1$ and $\langle\!\langle M_{12} S_1 \rangle\!\rangle = 2$, which implies $S_1 M_{12} \not\sim M_{12} S_1$.

Remark 9. The equality of class-change norms of any pair of branching bisimilar configurations $\alpha \sim \beta$ can help us to restrict the potential consistent responses $\beta = \beta_0 \xrightarrow{\tau} \beta_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k \xrightarrow{a} \beta'$ to a transition $\alpha \xrightarrow{a} \alpha'$. In particular, all β_i ($i \in \{0, 1, \ldots, k\}$) must have the same class-change norm. This is one of the points, which does not hold for weak bisimilarity.

Redundant variables characterize the suffixes. The fact of compositionality, i.e., the fact that $\sim$ is a congruence (Prop. 6), naturally leads us to look for possible decompositions, in particular for the "prime decompositions" of configurations α, as we do in Section 4. It will turn out that the decomposition of α might be different in the "suffix-context" $\alpha\gamma$ than in the context $\alpha\delta$ when $\gamma \not\sim \delta$. Nevertheless, the decompositions of α in these two contexts will turn out to be the same if $\textsc{red}(\gamma) = \textsc{red}(\delta)$, where for any $\beta \in \mathcal{V}^*$ we put

$$\textsc{red}(\beta) = \{X \in \mathcal{V} \mid X\beta \sim \beta\}.$$

The variables in $\textsc{red}(\beta)$ are called the *redundant variables w.r.t. β*. We observe:

Proposition 10. *We have $\alpha\beta \sim \beta$ if, and only if, $\alpha \in (\textsc{red}(\beta))^*$.*

Proof. The "if"-direction is obvious.

For the "only-if"-direction assume $\alpha_1 X \alpha_2 \beta \sim \beta$ where $\alpha_2 \in (\text{RED}(\beta))^*$ and $X \notin \text{RED}(\beta)$; hence $\alpha_1 X \alpha_2 \beta \sim \alpha_1 X \beta \sim \beta$ and $X\beta \not\sim \beta$. This implies that the cc-length of any path $X\beta \xrightarrow{u} \beta$ is positive, and thus $\langle\!\langle \alpha_1 X \beta \rangle\!\rangle \geq \langle\!\langle X\beta \rangle\!\rangle > \langle\!\langle \beta \rangle\!\rangle$, which excludes $\alpha_1 X \beta \sim \beta$ (since Prop. 8(2) implies $\langle\!\langle \alpha_1 X \beta \rangle\!\rangle = \langle\!\langle \beta \rangle\!\rangle$). □

To prepare a way for formalizing the above claims on decompositions, we first relativize the equivalence $\sim$ and the cc-norm $\langle\!\langle . \rangle\!\rangle$ w.r.t. the suffix-contexts.

For any $\gamma \in \mathcal{V}^*$ we define the relation $\sim_\gamma$ and the norm $\langle\!\langle . \rangle\!\rangle_\gamma$ as follows:

$$\alpha \sim_\gamma \beta \Leftrightarrow_{df} \alpha\gamma \sim \beta\gamma, \text{ and } \langle\!\langle \alpha \rangle\!\rangle_\gamma =_{df} \langle\!\langle \alpha\gamma \rangle\!\rangle - \langle\!\langle \gamma \rangle\!\rangle.$$

Remark. In Example 5 we have $\langle\!\langle C \rangle\!\rangle_\varepsilon = 3$, but $\langle\!\langle C \rangle\!\rangle_{M_2} = 2$ and $\langle\!\langle C \rangle\!\rangle_{M_{23}} = 1$. We can also check that $\text{RED}(S_i) = \emptyset$, $\text{RED}(M_R) = \{M_S \mid S \subseteq R\} \cup \{S_i \mid i \in R\}$, $\text{RED}(A) = \emptyset$, $\text{RED}(B) = \text{RED}(C) = \{S_1, M_1\}$, $\text{RED}(S_1 M_{12}) = \text{RED}(M_{12})$.

We note some simple facts:

Proposition 11.
(1) If $\gamma \sim \delta$, then $\sim_\gamma = \sim_\delta$, $\langle\!\langle . \rangle\!\rangle_\gamma = \langle\!\langle . \rangle\!\rangle_\delta$, and $\text{RED}(\gamma) = \text{RED}(\delta)$.
(2) If $\alpha \sim_\gamma \beta$, then $\langle\!\langle \alpha \rangle\!\rangle_\gamma = \langle\!\langle \beta \rangle\!\rangle_\gamma$.
(3) $\langle\!\langle \alpha\beta \rangle\!\rangle_\gamma = \langle\!\langle \alpha \rangle\!\rangle_{\beta\gamma} + \langle\!\langle \beta \rangle\!\rangle_\gamma$.
(4) $\alpha\gamma \sim \gamma$ iff $\alpha \in (\text{RED}(\gamma))^*$ iff $\langle\!\langle \alpha \rangle\!\rangle_\gamma = 0$;
in particular, $X \in \text{RED}(\gamma)$ iff $\langle\!\langle X \rangle\!\rangle_\gamma = 0$.

Proof. (1), (2), and (3) follow trivially from the definitions and the fact that $\alpha \sim_\gamma \beta$ implies $\langle\!\langle \alpha\gamma \rangle\!\rangle = \langle\!\langle \beta\gamma \rangle\!\rangle$ (by Prop. 8(2)).
(4) partly repeats Prop. 10, and otherwise it follows easily (by using (1), (2), (3)); in particular, if $\langle\!\langle \alpha \rangle\!\rangle_\gamma = \langle\!\langle \alpha\gamma \rangle\!\rangle - \langle\!\langle \gamma \rangle\!\rangle = 0$, then any witness $\alpha\gamma \xrightarrow{w} \varepsilon$ for $\alpha\gamma$ can be written $\alpha\gamma \xrightarrow{u} \gamma \xrightarrow{v} \varepsilon$ where $\alpha\gamma \xrightarrow{u} \gamma$ contains no class-change transition, which implies $\alpha\gamma \sim \gamma$. □

The next lemma says that any suffix-context is fully characterized by the set of respective redundant variables:

Lemma 12. *If $\text{RED}(\gamma) = \text{RED}(\delta)$, then $\sim_\gamma = \sim_\delta$ and $\langle\!\langle . \rangle\!\rangle_\gamma = \langle\!\langle . \rangle\!\rangle_\delta$.*

Proof. Suppose $\text{RED}(\gamma) = \text{RED}(\delta)$; we will verify that the set

$$\mathcal{B} = \{(\alpha\delta, \beta\delta) \mid \alpha\gamma \sim \beta\gamma\} \cup \sim$$

is a branching bisimulation. We note that $\mathcal{B}$ is symmetric, and consider

$$\text{a pair } (\mu, \nu) \in \mathcal{B} \text{ and a transition } \mu \xrightarrow{a} \mu'.$$

If $\mu \sim \nu$, then either $a = \tau$ and $\mu' \sim \nu$, in which case $(\mu', \nu) \in \mathcal{B}$, or there is a response $\nu = \nu_0 \xrightarrow{\tau} \nu_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \nu_k \xrightarrow{a} \nu'$ such that (μ', ν') and (μ, ν_i) for all $i \in \{1, 2, \ldots, k\}$ are in $\sim$, and thus in $\mathcal{B}$.

Now assume

$$(\mu,\nu) = (\alpha\delta, \beta\delta) \text{ where } \alpha\gamma \sim \beta\gamma.$$

If $\alpha = \varepsilon$, then $\gamma \sim \beta\gamma$, which entails $\beta \in (\text{RED}(\gamma))^* = (\text{RED}(\delta))^*$, and thus $\mu = \alpha\delta = \delta \sim \beta\delta = \nu$ (by Prop. 11); similarly, $\beta = \varepsilon$ also implies $\mu \sim \nu$. We thus assume that $\alpha \neq \varepsilon$ and $\beta \neq \varepsilon$; hence

$$\mu = \alpha\delta \xrightarrow{a} \alpha'\delta = \mu' \text{ for a transition } \alpha \xrightarrow{a} \alpha'. \tag{1}$$

Since $\alpha\gamma \sim \beta\gamma$, the transition $\alpha\gamma \xrightarrow{a} \alpha'\gamma$ entails that either $a = \tau$ and $\alpha'\gamma \sim \beta\gamma$, in which case $(\alpha'\delta, \beta\delta) = (\mu', \nu) \in \mathcal{B}$, or there is a corresponding response of one of the following forms:

1. $\beta\gamma = \beta_0\gamma \xrightarrow{\tau} \beta_1\gamma \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k\gamma \xrightarrow{a} \beta'\gamma$,
2. $\beta\gamma = \beta_0\gamma \xrightarrow{\tau} \beta_1\gamma \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k\gamma \xrightarrow{\tau} \gamma = \gamma_0 \xrightarrow{\tau} \gamma_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \gamma_\ell \xrightarrow{a} \gamma'$.

In the case 1 we have that $(\alpha'\gamma, \beta'\gamma)$ and $(\alpha\gamma, \beta_i\gamma)$ for all $i \in \{1, 2, \ldots, k\}$ are in $\sim$; hence $(\alpha'\delta, \beta'\delta)$ and $(\alpha\delta, \beta_i\delta)$ for all $i \in \{1, 2, \ldots, k\}$ are in $\mathcal{B}$, and thus $\nu = \beta\delta = \beta_0\delta \xrightarrow{\tau} \beta_1\delta \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k\delta \xrightarrow{a} \beta'\delta$ is an appropriate response to $\mu \xrightarrow{a} \mu'$ (i.e., to $\alpha\delta \xrightarrow{a} \alpha'\delta$ from (1)).

In the case 2 we have $\alpha\gamma \sim \gamma$, and thus also $\gamma \sim \beta\gamma$. Hence both α, β are from $(\text{RED}(\gamma))^* = (\text{RED}(\delta))^*$, which entails $\mu = \alpha\delta \sim \delta \sim \beta\delta = \nu$.

Since $\mathcal{B}$ is a branching bisimulation, we derive that $\sim_\gamma \subseteq \sim_\delta$, and by symmetry that $\sim_\gamma = \sim_\delta$. This also entails that $\langle\!\langle . \rangle\!\rangle_\gamma = \langle\!\langle . \rangle\!\rangle_\delta$. Indeed, we recall that $\langle\!\langle \alpha \rangle\!\rangle_\gamma$ is the smallest cc-length (counting only the class-changing transitions) of the paths $\alpha\gamma = \alpha_0\gamma \xrightarrow{a_1} \alpha_1\gamma \xrightarrow{a_2} \cdots \xrightarrow{a_n} \alpha_n\gamma = \gamma$. Since $\alpha_{i-1}\gamma \xrightarrow{a_i} \alpha_i\gamma$ is a class-changing transition (i.e., $\alpha_{i-1} \not\sim_\gamma \alpha_i$) iff $\alpha_{i-1}\delta \xrightarrow{a_i} \alpha_i\delta$ is a class-changing transition, we derive that $\langle\!\langle \alpha \rangle\!\rangle_\gamma = \langle\!\langle \alpha \rangle\!\rangle_\delta$. □

Remark 13. The fact that just $\text{RED}(\gamma)$ *determines whether* $\alpha\gamma \sim \beta\gamma$ *or* $\alpha\gamma \not\sim \beta\gamma$ *is not true in the case of weak bisimilarity, as illustrated by the example below.*

Example 14. *Consider the system defined by:*

$$\begin{array}{lll} A \xrightarrow{\tau} \varepsilon & B \xrightarrow{\tau} \varepsilon & X \xrightarrow{a} CX \\ A \xrightarrow{b} \varepsilon & B \xrightarrow{b} \varepsilon & X \xrightarrow{x} \varepsilon \\ A \xrightarrow{a} C & C \xrightarrow{c} \varepsilon & \end{array}$$

One can easily check that $\text{RED}(X) = \emptyset$, *clearly also* $\text{RED}(\varepsilon) = \emptyset$; *here we define the set of redundant variables with respect to the weak bisimilarity. However,* $AX \approx BX$, *while* $A \not\approx B$. *To see that* $A \not\approx B$ *note that* A *have a transition labelled by* a, *while* B *does not have such. Let us show now that indeed* $AX \approx BX$. *For every transition from* BX *there is an identical one from* AX. *The only transition of* AX, *which cannot be matched by identical one from* BX *is*

$$AX \xrightarrow{a} CX.$$

However the following response from BX *is completely fine*

$$BX \xrightarrow{\tau} X \xrightarrow{a} CX.$$

Note that the same response would not be correct in the case of branching bisimilarity as $X \not\sim AX$.

Relativization with respect to $R \subseteq \mathcal{V}$. We say that a *set* $R \subseteq \mathcal{V}$ is *suffix-generated* if $R = \mathrm{RED}(\gamma)$ for some $\gamma \in \mathcal{V}^*$; we often implicitly consider only suffix-generated $R \subseteq \mathcal{V}$ in what follows. For any (suffix-generated) set $R \subseteq \mathcal{V}$, Lemma 12 allows us to soundly define the *R-equivalence* $\sim_R$ by

$$\alpha \sim_R \beta \text{ if } \alpha\gamma \sim \beta\gamma \text{ for some } \gamma \text{ where } R = \mathrm{RED}(\gamma),$$

and the *cc_R-norm* $\langle\!\langle . \rangle\!\rangle_R$ by

$$\langle\!\langle \alpha \rangle\!\rangle_R = \langle\!\langle \alpha \rangle\!\rangle_\gamma \text{ for some } \gamma \text{ where } R = \mathrm{RED}(\gamma).$$

By the *cc_R-length* of a path $\alpha \xrightarrow{u} \beta$ we mean the number of transitions in the path that change the class of $\sim_R$. By

$$\text{an } R\text{-witness path for } \alpha$$

we mean a shortest path $\alpha \xrightarrow{u} \varepsilon$ whose cc_R-length is minimal, and thus equal to $\langle\!\langle \alpha \rangle\!\rangle_R$.

We also relativize RED (for redundant variables), putting

$$\mathrm{RED}_R(\alpha) = \mathrm{RED}(\alpha\gamma) \text{ for some } \gamma \text{ where } R = \mathrm{RED}(\gamma).$$

Abusing notation, we also write $\mathrm{RED}(\alpha, R)$ instead of $\mathrm{RED}_R(\alpha)$. We note that $\alpha \in R^*$ implies $\mathrm{RED}(\alpha, R) = R$.

The next proposition summarizes some consequences of the previous facts. We note that $\sim_R$ are not congruences in general, the respective properties (captured by (7)) are more subtle.

Proposition 15.
(1) $\mathrm{RED}(\varepsilon) = \emptyset$*;* $\sim = \sim_{\mathrm{RED}(\varepsilon)}$*;* $\langle\!\langle . \rangle\!\rangle = \langle\!\langle . \rangle\!\rangle_{\mathrm{RED}(\varepsilon)}$*.*
(2) $\langle\!\langle \alpha \rangle\!\rangle_R \leq \|\alpha\|$*.*
(3) $\alpha \sim_R \beta$ *implies* $\langle\!\langle \alpha \rangle\!\rangle_R = \langle\!\langle \beta \rangle\!\rangle_R$*.*
(4) $\langle\!\langle \alpha\beta \rangle\!\rangle_R = \langle\!\langle \alpha \rangle\!\rangle_{R'} + \langle\!\langle \beta \rangle\!\rangle_R$ *where* $R' = \mathrm{RED}(\beta, R)$*.*
(5) $\alpha \sim_R \varepsilon$ *iff* $\alpha \in R^*$ *iff* $\langle\!\langle \alpha \rangle\!\rangle_R = 0$*; in particular,* $\langle\!\langle X \rangle\!\rangle_R = 0$ *iff* $X \in R$*.*
(6) If $\alpha \sim_R \beta$*, then* $\mathrm{RED}(\alpha, R) = \mathrm{RED}(\beta, R)$*.*
(7) Suppose $\alpha \sim_R \beta$ *and* $R' = \mathrm{RED}(\alpha, R)$ *(*$= \mathrm{RED}(\beta, R)$*).*
Then $\gamma \sim_{R'} \delta$ *iff* $\gamma\alpha \sim_R \delta\beta$ *(for any* γ, δ*).*

Proof. Points (1)–(6) are routine consequences of the definitions and previous facts; in particular, $\mathrm{RED}(\varepsilon) = \emptyset$ since we have excluded silent variables.

(7): Let $\alpha \sim_R \beta$ and $R' = \mathrm{RED}(\alpha, R)$. We fix μ such that $\alpha\mu \sim \beta\mu$ and $\mathrm{RED}(\mu) = R$; hence $\mathrm{RED}(\alpha\mu) = \mathrm{RED}(\alpha, R) = R'$.

We thus get: $\gamma \sim_{R'} \delta$ iff $\gamma\alpha\mu \sim \delta\alpha\mu$ iff $\gamma\alpha\mu \sim \delta\beta\mu$ iff $\gamma\alpha \sim_R \delta\beta$. □

4 Relative prime decomposition

Recall that we implicitly assume a normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$.

Redundancy-free form, relative prime variables, relative prime form. We say that $\alpha \in \mathcal{V}^*$ is *R-redundancy-free*, for $R \subseteq \mathcal{V}$, if we do not have $\alpha = \beta X \gamma$ where $X \in \textsc{red}(\gamma, R)$. In other words, ε is R-redundancy-free for any R, and αX is R-redundancy-free if $X \notin R$ and α is R'-redundancy-free for $R' = \textsc{red}(X, R)$. When saying *redundancy-free*, we mean R-redundancy-free for $R = \textsc{red}(\varepsilon) = \emptyset$.

We define the *R-redundancy-free form* $\textsc{rff}_R(\alpha)$ of $\alpha \in \mathcal{V}^*$ inductively:

i) $\textsc{rff}_R(\varepsilon) = \varepsilon$;
ii) if $X \in R$, then $\textsc{rff}_R(\beta X) = \textsc{rff}_R(\beta)$;
iii) if $X \notin R$, then $\textsc{rff}_R(\beta X) = \textsc{rff}_{R'}(\beta)\, X$ where $R' = \textsc{red}(X, R)$.

We note that $\alpha \sim_R \textsc{rff}_R(\alpha)$, and α is R-redundancy-free iff $\alpha = \textsc{rff}_R(\alpha)$.

Remark. Technically we could omit the explicit definition of $\textsc{rff}_R(\alpha)$. It serves us mainly as a demonstration that our inductive definitions and proofs for configurations proceed in the "right-to-left" (or "bottom-up") fashion.

A *variable* $A \notin R$ is *R-decomposable* if $A \sim_R \alpha$ where α is R-redundancy-free and $|\alpha| > 1$; if $A \notin R$ is not R-decomposable, then A is *non-R-decomposable*.

Proposition 16.
(1) If α is R-redundancy-free, then $\langle\!\langle \alpha \rangle\!\rangle_R \geq |\alpha|$.
(2) If A is R-decomposable, then $A \sim_R \beta B$ where B is non-R-decomposable and $\langle\!\langle A \rangle\!\rangle_R > \langle\!\langle B \rangle\!\rangle_R \geq 1$.

Proof. (1) follows from the already observed fact that $X\gamma \not\sim \gamma$ implies that the cc-length of any path $X\gamma \xrightarrow{u} \gamma$ is positive.

(2): If A is R-decomposable, then by definition we have $A \sim_R \beta B$ where βB is R-redundancy-free, $\beta \neq \varepsilon$, and $\langle\!\langle A \rangle\!\rangle_R = \langle\!\langle \beta \rangle\!\rangle_{R'} + \langle\!\langle B \rangle\!\rangle_R$ where $R' = \textsc{red}(B, R)$ (which follows from the definitions and Prop. 15). Since $B \notin R$ and β is R'-redundancy free, we have $\langle\!\langle A \rangle\!\rangle_R > \langle\!\langle B \rangle\!\rangle_R \geq 1$. Similarly, if B is R-decomposable, then $B \sim_R \delta C$ where $\langle\!\langle B \rangle\!\rangle_R > \langle\!\langle C \rangle\!\rangle_R \geq 1$, which entails $A \sim_R \beta B \sim_R \beta\delta C$. We could thus immediately take β and B so that $\langle\!\langle B \rangle\!\rangle_R$ is minimal, in which case B is non-R-decomposable. □

***R*-prime form.** For any (suffix-generated) $R \subseteq \mathcal{V}$, the equivalence $\sim_R$ induces a partition on the set of all non-R-decomposable variables (from $\mathcal{V} \setminus R$). In each class of this partition (containing non-R-decomposable variables that are pairwise equivalent w.r.t. $\sim_R$) we choose a variable and call it an *R-prime*.

Convention. We further assume that the R-primes for all (suffix-generated) sets $R \subseteq \mathcal{V}$ have been fixed, in our assumed nBPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, unless stated otherwise.

We say that $\alpha \in \mathcal{V}^*$ is in the *R-prime form* if

– either $\alpha = \varepsilon$,
– or $\alpha = \beta X$ where X is an R-prime and β is in the R'-prime form for $R' = \textsc{red}(X, R)$.

The next lemma shows that for any α and R there is the unique α' in the R-prime form such that $\alpha \sim_R \alpha'$.

Lemma 17.
(1) For any $\alpha \in \mathcal{V}^*$ and $R \subseteq \mathcal{V}$ there is α' in the R-prime form where $\alpha \sim_R \alpha'$.
(2) If $\alpha \sim_R \beta$ and both α and β are in the R-prime form, then $\alpha = \beta$.

Proof. (1): We show the claim for α, R by induction on $\langle\!\langle \alpha \rangle\!\rangle_R$. If $\langle\!\langle \alpha \rangle\!\rangle_R = 0$, i.e. $\alpha \in R^*$, then $\alpha \sim_R \varepsilon$ (where ε is in the R-prime form). We thus assume $\alpha = \beta X \gamma$ where $\gamma \in R^*$ and $X \notin R$; note that $\alpha \sim_R \beta X \gamma \sim_R \beta X$. Recalling Prop. 16(2), we deduce that $X \sim_R \delta Y$ where Y is non-R-decomposable (maybe $\delta = \varepsilon$ and $X = Y$); hence $\alpha \sim_R \beta X \sim_R \beta\delta Y \sim_R \beta\delta Z$ where Z is the (unique) R-prime such that $Y \sim_R Z$. We put $R' = \textsc{red}(Z, R)$, and note that $\langle\!\langle \alpha \rangle\!\rangle_R = \langle\!\langle \beta\delta \rangle\!\rangle_{R'} + \langle\!\langle Z \rangle\!\rangle_R$; hence $\langle\!\langle \beta\delta \rangle\!\rangle_{R'} < \langle\!\langle \alpha \rangle\!\rangle_R$. By the induction hypothesis there is β' in the R'-prime form such that $\beta' \sim_{R'} \beta\delta$; hence $\alpha \sim_R \beta' Z$ and $\beta' Z$ is in the R-prime form.

(2): Suppose the claim is not true. Then there are $\alpha \sim_R \beta$, both α and β being in the R-prime form, such that $\alpha = \alpha' X \gamma$ and $\beta = \beta' Y \gamma$ where X, Y are different R'-primes for $R' = \textsc{red}(\gamma, R)$; hence $\alpha' X \sim_{R'} \beta' Y$, and by symmetry we can assume $\langle\!\langle X \rangle\!\rangle_{R'} \geq \langle\!\langle Y \rangle\!\rangle_{R'}$.

More generally, if the claim is not true, then we deduce that there are some R, α, β, and two *different* R-primes X, Y where

$$\alpha X \sim_R \beta Y \text{ and } \langle\!\langle X \rangle\!\rangle_R \geq \langle\!\langle Y \rangle\!\rangle_R.$$

We take such a counterexample with a shortest possible R-witness path for αX. (Note that we do not require that αX and βY are in the R-prime form.)

We note that $\alpha X \not\sim_R X$ (since otherwise $X \sim_R \beta Y$, and thus either $\langle\!\langle X \rangle\!\rangle_R > \langle\!\langle Y \rangle\!\rangle_R$ and X is R-decomposable, or $X \sim_R \beta Y \sim_R Y$, which is impossible by our definition of R-primes). We thus also have $\alpha X \not\sim_R Y$ (since $\langle\!\langle \alpha X \rangle\!\rangle_R > \langle\!\langle Y \rangle\!\rangle_R$).

Hence the first step in the R-witness path for αX is $\alpha X \xrightarrow{a} \alpha' X$ (where $\alpha' X$ has a shorter R-witness path than αX). Since $\alpha X \sim_R \beta Y$, and $\alpha X \not\sim_R Y$, this first step must have a response $\beta Y = \beta_0 Y \xrightarrow{\tau} \beta_1 Y \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k Y \xrightarrow{a} \beta' Y$ where $\alpha' X \sim_R \beta' Y$ (and all β_i are nonempty since $\alpha X \not\sim_R Y$); this contradicts the choice of our counterexample. □

Lemma 17 allows us to soundly define

$\mathrm{PD}_R(\alpha)$, the *prime decomposition of* α w.r.t. (i.e., in the context of) R,

as the unique configuration in the R-prime form such that $\alpha \sim_R \mathrm{PD}_R(\alpha)$.

We explicitly note the following consequences:

Corollary 18.
(1) For any $R \subseteq \mathcal{V}$ and $\alpha, \beta \in \mathcal{V}^*$ we have $\alpha \sim_R \beta$ iff $\mathrm{PD}_R(\alpha) = \mathrm{PD}_R(\beta)$.
(2) The mapping $\mathrm{PD}_R(\alpha)$ (with arguments R, α) satisfies the following:

i) $\mathrm{PD}_R(\varepsilon) = \varepsilon$;
ii) if $X \in R$, then $\mathrm{PD}_R(X) = \varepsilon$;
iii) if X is an R-prime, then $\mathrm{PD}_R(X) = X$;
iv) if $X \notin R$, then $\mathrm{PD}_R(X) = \mathrm{PD}_R(\beta Z)$ *for the unique R-prime Z and any β such that $X \sim_R \beta Z$;*

v) $\mathrm{PD}_R(\alpha X) = \mathrm{PD}_{R'}(\alpha)\,\mathrm{PD}_R(X)$ *where* $R' = \mathrm{RED}(\mathrm{PD}_R(X), R)$.

Remark. In the Example 5, M_{12} is non-$\emptyset$-decomposable, while A is $\emptyset$-decomposable since $A \sim_\emptyset S_1M_3$. We have $\mathrm{PD}_\emptyset(C) = M_1M_3M_2$, so $\mathrm{PD}_{\{M_2,S_2\}}(C) = M_1M_3$ and $\mathrm{PD}_{\{M_{23},M_2,M_3,S_2,S_3\}}(C) = M_1$. Therefore $\mathrm{PD}_\emptyset(S_2CM_{23}) = S_2M_1M_{23}$.

We note that we can have, e.g., $\mathrm{PD}_R(X) = XY$, which may seem a bit unnatural. This is illustrated by the example given by the following rules:

$$X \xrightarrow{a} Y,\ Y \xrightarrow{b} Y,\ Y \xrightarrow{\tau} \varepsilon.$$

Here Y is $\emptyset$-prime, with $\langle\!\langle Y \rangle\!\rangle_\emptyset = 1$, but X is not $\emptyset$-prime, since $X \sim_\emptyset XY$ and $\langle\!\langle X \rangle\!\rangle_\emptyset = 2$. We have $Y \in \mathrm{RED}(Y, \emptyset)$ (since $Y \sim_\emptyset YY$), and X is $\{Y\}$-prime, with $\langle\!\langle X \rangle\!\rangle_{\{Y\}} = 1$.

5 Consistent bases

In this section we show that R-primes and R-decompositions of non-R-primes of exponentially bounded sizes can be guessed, and the consistency of the guess can be verified. The candidates for decompositions will be captured by so called *bases*; these are defined via a technical notion of *pre-bases*.

We again assume a given normed BPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, now with no information about R-primes. We now use the symbol $\mathcal{B}$ for (pre)bases, rather than for bisimulations.

Pre-bases. A *pre-base* $\mathcal{B}$ is determined by its domain $\mathrm{DOM}(\mathcal{B}) \subseteq 2^\mathcal{V}$, where $\emptyset \in \mathrm{DOM}(\mathcal{B})$, and two disjoint sets $\mathcal{B}_{dec}$ (*decompositions*) and $\mathcal{B}_{prop}$ (*propagations*) containing triples (A, α, R) where $A \in \mathcal{V}$, $\alpha \in \mathcal{V}^*$, $R \subseteq \mathcal{V}$, for which the following conditions hold:

1. For each $R \in \mathrm{DOM}(\mathcal{B})$, the set $\mathcal{V} \setminus R$ is partitioned into the set of $(\mathcal{B}, R)$-*primes* and the set of $(\mathcal{B}, R)$-*non-primes*. (Thus for $R = \emptyset$ each $X \in \mathcal{V}$ is either a $(\mathcal{B}, R)$-prime or a $(\mathcal{B}, R)$-non-prime.) Moreover:
   - For each $(\mathcal{B}, R)$-non-prime A there is precisely one triple (A, α, R) in $\mathcal{B}_{dec}$, where $\alpha = \beta B$ for a $(\mathcal{B}, R)$-prime B (and β can be empty).
   - For each $(\mathcal{B}, R)$-prime B there might be some triples (B, XB, R) in $\mathcal{B}_{prop}$ (where $X \in \mathcal{V}$). We put
   $$\mathrm{RED}^\mathcal{B}(B, R) = \{X \mid (B, XB, R) \in \mathcal{B}_{prop}\}. \tag{2}$$
2. $\mathcal{B}_{dec}$ and $\mathcal{B}_{prop}$ do not contain any other triples than those mentioned above.
3. $\mathrm{DOM}(\mathcal{B})$ is equal to the least set that contains $\emptyset$ and is closed under propagation, which means that if $R' = \mathrm{RED}^\mathcal{B}(B, R)$ for some $R \in \mathrm{DOM}(\mathcal{B})$ and some $(\mathcal{B}, R)$-prime B, then $R' \in \mathrm{DOM}(\mathcal{B})$.

Informally, a triple $(A, \alpha, R) \in \mathcal{B}_{dec}$ can be viewed as a statement that A is not an R-prime and that $\mathrm{PD}_R(A) = \alpha$ (which includes the case $\alpha = B$ for a guessed R-prime B). A triple $(A, XA, R) \in \mathcal{B}_{prop}$ can be viewed as a statement that $X \in \textsc{red}(A, R)$ (and that A is an R-prime).

$(\mathcal{B}, R)$-prime-decompositions. For a pre-base $\mathcal{B}$ and $R \in \textsc{dom}(\mathcal{B})$ we define the *$(\mathcal{B}, R)$-prime-decomposition form* $\mathrm{PD}^{\mathcal{B}}_R(\gamma)$ of strings $\gamma \in \mathcal{V}^*$ by the following (inductive) definition:

i. $\mathrm{PD}^{\mathcal{B}}_R(\varepsilon) = \varepsilon$;
ii. if $A \in R$, then $\mathrm{PD}^{\mathcal{B}}_R(\beta A) = \mathrm{PD}^{\mathcal{B}}_R(\beta)$;
iii. if A is a $(\mathcal{B}, R)$-prime, then $\mathrm{PD}^{\mathcal{B}}_R(\beta A) = \mathrm{PD}^{\mathcal{B}}_{R'}(\beta)\, A$ for $R' = \textsc{red}^{\mathcal{B}}(A, R)$;
iv. if A is a $(\mathcal{B}, R)$-non-prime, then $\mathrm{PD}^{\mathcal{B}}_R(\beta A) = \mathrm{PD}^{\mathcal{B}}_R(\beta\alpha)$ where $(A, \alpha, R) \in \mathcal{B}_{dec}$.

Our definition of pre-bases has not excluded that $\mathrm{PD}^{\mathcal{B}}_R(\gamma)$ could be infinite. Nevertheless, our intention is that a triple $(A, \alpha, R) \in \mathcal{B}_{dec}$ captures the guess $\mathrm{PD}_R(A) = \alpha$, which entails $|\alpha| \leq \langle\!\langle \alpha \rangle\!\rangle_R = \langle\!\langle A \rangle\!\rangle_R \leq \|A\|$. We thus impose further (syntactic) conditions:

Bases. A base $\mathcal{B}$ is a pre-base that satisfies:

1. For all $A \in \mathcal{V}$ and $R \in \textsc{dom}(\mathcal{B})$ we have $|\mathrm{PD}^{\mathcal{B}}_R(A)| \leq \|A\|$.
2. For each $(A, \alpha, R) \in \mathcal{B}_{dec}$ we have $\alpha = \mathrm{PD}^{\mathcal{B}}_R(\alpha)$.

It is easy to verify the next fact:

Proposition 19. *The size of any base is exponentially bounded, and the respective conditions can be verified in exponential-time (w.r.t. the size of $\mathcal{G}$).*

Base-generated equivalences. For a base $\mathcal{B}$ and any $R \in \textsc{dom}(\mathcal{B})$ we define the relation $\equiv^{\mathcal{B}}_R$ as follows:

$$\alpha \equiv^{\mathcal{B}}_R \beta \Leftrightarrow_{df} \mathrm{PD}^{\mathcal{B}}_R(\alpha) = \mathrm{PD}^{\mathcal{B}}_R(\beta).$$

By definition, $A \equiv^{\mathcal{B}}_R \alpha$ for any triple (A, α, R) in $\mathcal{B}$ (i.e., for any $(A, \alpha, R) \in \mathcal{B}_{dec}$ and any $(A, XA, R) \in \mathcal{B}_{prop}$; in the latter case we have $X \in \textsc{red}^{\mathcal{B}}(A, R)$).

We put

$$\mathrm{PD}^{\mathcal{B}}(\alpha) = \mathrm{PD}^{\mathcal{B}}_{\emptyset}(\alpha) \text{ and } \equiv^{\mathcal{B}} = \equiv^{\mathcal{B}}_{\emptyset}.$$

We also generalize (2), defining $\textsc{red}^{\mathcal{B}}_R(\alpha)$, rather written as $\textsc{red}^{\mathcal{B}}(\alpha, R)$, by the following inductive definition:

- $\textsc{red}^{\mathcal{B}}(\varepsilon, R) = R$ (in particular, $\textsc{red}^{\mathcal{B}}(\varepsilon, \emptyset) = \emptyset$);
- if B is a $(\mathcal{B}, R)$-prime, then $\textsc{red}^{\mathcal{B}}(\alpha B, R) = \textsc{red}^{\mathcal{B}}(\alpha, R')$ where $R' = \textsc{red}^{\mathcal{B}}(B, R)$;
- $\textsc{red}^{\mathcal{B}}(\alpha, R) = \textsc{red}^{\mathcal{B}}(\mathrm{PD}^{\mathcal{B}}_R(\alpha), R)$.

Intended base. Given an nBPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, an *intended base* $\mathcal{B}$ arises by choosing the R-primes as described before Lemma 17, putting $(A, \mathrm{PD}_R(\alpha), R)$ in $\mathcal{B}_{dec}$ for each suffix-generated R and each non-R-prime A, and putting (A, XA, R) in $\mathcal{B}_{prop}$ iff A is an R-prime and $X \in \mathrm{RED}(A, R)$.

Corollary 18(1) shows that for any intended base $\mathcal{B}$ we have $\equiv^{\mathcal{B}}_R = \sim_R$; in particular, $\equiv^{\mathcal{B}} = \sim$, i.e., $\equiv^{\mathcal{B}}$ is the maximal branching bisimulation.

Now we look for some suitable "syntactic" conditions guaranteeing that $\equiv^{\mathcal{B}}$, for a given (not necessarily intended) base $\mathcal{B}$, is a branching bisimulation.

Each $\equiv^{\mathcal{B}}$ is symmetric, and we need to guarantee that for each $\alpha \equiv^{\mathcal{B}} \beta$ and each transition $\alpha \xrightarrow{a} \alpha'$ we have that either $a = \tau$ and $\alpha' \equiv^{\mathcal{B}} \beta$, or there is a corresponding response $\beta = \beta_0 \xrightarrow{\tau} \beta_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k \xrightarrow{a} \beta'$ where $\alpha' \equiv^{\mathcal{B}} \beta'$ and $\alpha \equiv^{\mathcal{B}} \beta_i$ for all $i \in \{1, 2, \ldots, k\}$.

Legal-move outcomes, consistent bases. A natural idea is to define *the set of $(\mathcal{B}, R)$-legal (move) outcomes for $\alpha \in \mathcal{V}^*$*, denoted $\mathrm{LO}^{\mathcal{B}}_R(\alpha)$ as a subset of $\mathcal{A} \times \mathcal{V}^*$:

$(a, \alpha') \in \mathrm{LO}^{\mathcal{B}}_R(\alpha)$ if

i) either $a = \tau$ and $\alpha' = \mathrm{PD}^{\mathcal{B}}_R(\alpha)$,
ii) or there is a sequence $\alpha = \alpha_0 \xrightarrow{\tau} \alpha_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \alpha_k \xrightarrow{a} \alpha''$ where $\alpha' = \mathrm{PD}^{\mathcal{B}}_R(\alpha'')$ and $\mathrm{PD}^{\mathcal{B}}_R(\alpha_i) = \mathrm{PD}^{\mathcal{B}}_R(\alpha)$ for all $i \in \{1, 2, \ldots, k\}$.

A *pair* (α, β) is $(\mathcal{B}, R)$-*consistent* if $\alpha \equiv^{\mathcal{B}}_R \beta$ and $\mathrm{LO}^{\mathcal{B}}_R(\alpha) = \mathrm{LO}^{\mathcal{B}}_R(\beta)$.

A *base* $\mathcal{B}$ is *consistent* if for each (A, α, R) in $\mathcal{B}$ we have that (A, α) is a $(\mathcal{B}, R)$-consistent pair.

Rest of a proof of Theorem 2. We now show that checking consistency of a base can be done in exponential-time w.r.t. the size of $\mathcal{G}$, that any intended base is consistent, and that $\equiv^{\mathcal{B}}$ is a branching bisimulation for any consistent base $\mathcal{B}$. After establishing these facts, a desired algorithm is obvious.

Lemma 20. *Checking consistency of a base can be done in exponential time (w.r.t. the size of $\mathcal{G}$).*

Proof. Given a base $\mathcal{B}$ (whose size is exponentially bounded w.r.t. $\mathcal{G}$ by definition), we need to verify that $\mathrm{LO}^{\mathcal{B}}_R(A) = \mathrm{LO}^{B}_R(\alpha)$ for each triple (A, α, R) in $\mathcal{B}$. There are at most exponentially many triples, so it suffices to focus on one of them; we have either $(A, \alpha, R) \in \mathcal{B}_{dec}$, in which case $\alpha = \mathrm{PD}^{\mathcal{B}}_R(A)$, or $(A, \alpha, R) = (A, XA, R) \in \mathcal{B}_{prop}$; in both cases we thus have $\mathrm{PD}^{\mathcal{B}}_R(A) = \mathrm{PD}^{\mathcal{B}}_R(\alpha)$, i.e., $A \equiv^{\mathcal{B}}_R \alpha$.

It suffices to show how to compute $\mathrm{LO}^{\mathcal{B}}_R(\beta)$ (for any β). We first define the *τ-closure* $\mathrm{TC}^{\mathcal{B}}_R(\beta)$ as the least set $\mathcal{T}$ that contains β and satisfies the following condition:

for any $Y\gamma \in \mathcal{T}$, $R' = \mathrm{RED}^{\mathcal{B}}_R(\gamma, R)$, and a rule $Y \xrightarrow{\tau} \delta$ of $\mathcal{G}$, we have:

i) if $Y \in R'$ and $\delta \in (R')^*$, then $\delta\gamma \in \mathcal{T}$;
ii) if $Y \notin R'$ and $\delta = \delta_1 Z \delta_2$ where $\delta_2 \in (R')^*$, $Z \notin R'$, and $\mathrm{PD}^{\mathcal{B}}_R(\delta_1 Z) = \mathrm{PD}^{\mathcal{B}}_R(Y)$, then $\delta_1 Z\gamma \in \mathcal{T}$.

We can easily check that $\mathrm{LO}^{\mathcal{B}}_R(\beta) = \{(\tau, \mathrm{PD}^{\mathcal{B}}_R(\beta))\} \cup \{(a, \mathrm{PD}^{\mathcal{B}}_R(\mu)) \mid \gamma \xrightarrow{a} \mu \text{ for some } \gamma \in \mathrm{TC}^{\mathcal{B}}_R(\beta)\}$. A dynamic-programming algorithm computing $\mathrm{LO}^{\mathcal{B}}_R(\beta)$ is thus obvious; its time-complexity is bounded by the length of β multiplied by an exponential function of the size of $\mathcal{G}$. □

Lemma 21. *Any intended base is consistent.*

Proof. Suppose $\mathcal{B}$ is an intended base. By the definition of bases, any (A, α, R) in $\mathcal{B}$ satisfies $A \equiv^{\mathcal{B}}_R \alpha$; moreover, we have $PD^{\mathcal{B}}_R(A) = PD^{\mathcal{B}}_R(\alpha) \neq \varepsilon$. We also note that $(\tau, PD^{\mathcal{B}}_R(A))$ belongs to both $\mathrm{LO}^{\mathcal{B}}_R(A)$ and $\mathrm{LO}^{\mathcal{B}}_R(\alpha)$.

Let $(a, \beta') \in \mathrm{LO}^{\mathcal{B}}_R(A)$, where $a \neq \tau$ or $\beta' \neq \mathrm{PD}^{\mathcal{B}}_R(A)$, due to the sequence

$$A = \beta_0 \xrightarrow{\tau} \beta_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \beta_k \xrightarrow{a} \beta''$$

where $\mathrm{PD}^{\mathcal{B}}_R(\beta_i) = \mathrm{PD}^{\mathcal{B}}_R(A)$ for all $i \in \{0, 1, \ldots, k\}$ and $\beta' = \mathrm{PD}^{\mathcal{B}}_R(\beta'')$. Since $\mathcal{B}$ is an intended base (and thus $\equiv^{\mathcal{B}}_R = \sim_R$), we have $A \sim_R \alpha \sim_R \beta_0 \sim_R \beta_1 \sim_R \cdots \sim_R \beta_k$; we also have either $a \neq \tau$ or $\beta'' \not\sim_R A$. There must be a response from α (composed from responses to the transitions $\beta_{i-1} \xrightarrow{\tau} \beta_i$ and $\beta_k \xrightarrow{a} \beta''$) of the form

$$\alpha = \alpha_0 \xrightarrow{\tau} \alpha_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \alpha_\ell \xrightarrow{a} \alpha''$$

where $\alpha_i \sim_R A$ for all $i \in \{0, 1, \ldots, \ell\}$, and $\alpha'' \sim_R \beta''$. Hence $PD^{\mathcal{B}}_R(\alpha_i) = PD^{\mathcal{B}}_R(A)$ (for all $i \in \{0, 1, \ldots, \ell\}$) and $PD^{\mathcal{B}}_R(\alpha'') = PD^{\mathcal{B}}_R(\beta'') = \beta'$. Therefore $(a, \beta') \in \mathrm{LO}^{\mathcal{B}}_R(\alpha)$, and thus $\mathrm{LO}^{\mathcal{B}}_R(A) \subseteq \mathrm{LO}^{\mathcal{B}}_R(\alpha)$.

Analogously we derive $\mathrm{LO}^{\mathcal{B}}_R(\alpha) \subseteq \mathrm{LO}^{\mathcal{B}}_R(A)$, and thus $\mathrm{LO}^{\mathcal{B}}_R(A) = \mathrm{LO}^{\mathcal{B}}_R(\alpha)$. □

The remaining fact is captured by Lemma 23; its main technical point is shown by the next proposition:

Proposition 22. *If $\mathcal{B}$ is a consistent base and $\alpha, \beta \in \mathcal{V}^*$ satisfy $\mathrm{PD}^{\mathcal{B}}_R(\alpha) = \mathrm{PD}^{\mathcal{B}}_R(\beta) \neq \varepsilon$, then $\mathrm{LO}^{\mathcal{B}}_R(\alpha) = \mathrm{LO}^{\mathcal{B}}_R(\beta)$.*

Proof. Suppose $\mathcal{B}$ is a consistent base and $\mathrm{PD}^{\mathcal{B}}_R(\alpha) = \mathrm{PD}^{\mathcal{B}}_R(\beta) \neq \varepsilon$. The definition of $\mathrm{PD}^{\mathcal{B}}_R(..)$ implies that α can be "transformed" into β by using triples (A, γ, R') from $\mathcal{B}$; we first transform α into $\mathrm{PD}^{\mathcal{B}}_R(\alpha)$, and then, by a "backward-transformation" we change $\mathrm{PD}^{\mathcal{B}}_R(\alpha) = \mathrm{PD}^{\mathcal{B}}_R(\beta)$ into β. There is thus a (transformation) sequence

$$\alpha = \gamma_0, \gamma_1, \ldots, \gamma_m = \beta$$

where $\mathrm{PD}^{\mathcal{B}}_R(\gamma_i) = \mathrm{PD}^{\mathcal{B}}_R(\alpha)$ for all $i \in \{0, 1, \ldots, m\}$ and each pair (γ_i, γ_{i+1}) satisfies the following:

one of γ_i, γ_{i+1} is in the form $\nu X \mu$ and the other in the form $\nu \delta \mu$,

where $\mu = \mathrm{PD}^{\mathcal{B}}_R(\mu)$ and for $R' = \textsc{red}^{\mathcal{B}}(\mu, R)$ we have

i) either $X \notin R'$ and $(X, \delta, R') \in \mathcal{B}_{dec}$,
ii) or $X \in R'$ and $\delta = \varepsilon$.

It thus suffices to check that we have $\mathrm{LO}^{\mathcal{B}}_R(\gamma_i) = \mathrm{LO}^{\mathcal{B}}_R(\gamma_{i+1})$ for each $i \in \{0, 1, \ldots, m-1\}$. For a given pair (γ_i, γ_{i+1}) in the above form we now aim to show that

$$\mathrm{LO}^{\mathcal{B}}_R(\nu X\mu) = \mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu).$$

By our definitions, $(\tau, \mathrm{PD}^{\mathcal{B}}_R(\nu X\mu))$ is in both $\mathrm{LO}^{\mathcal{B}}_R(\nu X\mu)$ and $\mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu)$. Suppose now that $(a, \rho) \in \mathrm{LO}^{\mathcal{B}}_R(\nu X\mu)$, where $a \neq \tau$ or $\rho \neq \mathrm{PD}^{\mathcal{B}}_R(\nu X\mu)$, due to a sequence

$$\nu X\mu = \rho_0 \xrightarrow{\tau} \rho_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \rho_k \xrightarrow{a} \rho' \tag{3}$$

where $\mathrm{PD}^{\mathcal{B}}_R(\rho_i) = \mathrm{PD}^{\mathcal{B}}_R(\nu X\mu)$ (for all $i \in \{0, 1, \ldots, k\}$) and $\mathrm{PD}^{\mathcal{B}}_R(\rho') = \rho$. If (3) "does not use X", i.e., can be written as

$$\nu X\mu = \nu_0 X\mu \xrightarrow{\tau} \nu_1 X\mu \xrightarrow{\tau} \cdots \xrightarrow{\tau} \nu_k X\mu \xrightarrow{a} \nu' X\mu \tag{4}$$

(where $\nu = \nu_0 \xrightarrow{\tau} \nu_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} \nu_k \xrightarrow{a} \nu'$), then we obviously have $(a, \rho) \in \mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu)$ (as shown by (4) when X is replaced with δ).

We recall that $R' = \textsc{red}^{\mathcal{B}}(\mu, R)$, and put $R'' = \textsc{red}^{\mathcal{B}}(X, R') = \textsc{red}^{\mathcal{B}}(\delta, R')$.

If $\nu \notin (R'')^*$, then (3) must be always of the form (4), by definition of legal outcomes; we thus further assume $\nu \in (R'')^*$. If $X \notin R'$ (and thus $\delta \notin (R')^*$), then (3) is either of the form (4) or it uses X but not μ; in the latter case we have a sequence $\nu X\mu = \nu_0 X\mu \xrightarrow{\tau} \nu_1 X\mu \xrightarrow{\tau} \cdots \xrightarrow{\tau} \nu_k X\mu \cdots \xrightarrow{a} \bar{\rho}\mu$ where $\nu_k = \varepsilon$ and $(a, \rho) = (a, \bar{\rho}\mu)$ where $(a, \bar{\rho})$ is in $\mathrm{LO}^{\mathcal{B}}_{R'}(X)$, and thus also in $\mathrm{LO}^{\mathcal{B}}_{R'}(\delta)$ since $(X, \delta, R') \in \mathcal{B}_{dec}$ and $\mathcal{B}$ is consistent; hence $(a, \bar{\rho}\mu) = (a, \rho) \in \mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu)$.

If ($\nu \in (R'')^*$ and) $X \in R'$, then $\delta = \varepsilon$ and $\mu \neq \varepsilon$ (since $\mathrm{PD}^{\mathcal{B}}_R(\alpha) \neq \varepsilon$); let us write $\mu = A\mu'$ where A is a $(\mathcal{B}, R''')$-prime for $R''' = \textsc{red}(\mu', R)$ and we have $(A, XA, R''') \in \mathcal{B}_{prop}$. Now (3) can also "use A", but cannot "use μ'", in which case $(a, \rho) = (a, \bar{\rho}\,\mu')$ where $(a, \bar{\rho})$ is in $\mathrm{LO}^{\mathcal{B}}_{R'''}(XA)$. Due to consistency of $\mathcal{B}$ we have $(a, \bar{\rho}) \in \mathrm{LO}^{\mathcal{B}}_{R'''}(A)$, and thus $(a, \bar{\rho}\mu') = (a, \rho) \in \mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu)$.

Hence $\mathrm{LO}^{\mathcal{B}}_R(\nu X\mu) \subseteq \mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu)$; the direction $\mathrm{LO}^{\mathcal{B}}_R(\nu\delta\mu) \subseteq \mathrm{LO}^{\mathcal{B}}_R(\nu X\mu)$ follows analogously. □

Lemma 23. *If $\mathcal{B}$ is a consistent base, then $\equiv^{\mathcal{B}}$ is a branching bisimulation.*

Proof. Suppose $\mathcal{B}$ is a consistent base, but $\equiv_{\mathcal{B}}$ is not a branching bisimulation. Then there is a pair (α, β), where $\mathrm{PD}^{\mathcal{B}}(\alpha) = \mathrm{PD}^{\mathcal{B}}(\beta)$, and a transition $\alpha \xrightarrow{a} \alpha'$ that has no adequate response (w.r.t. $\equiv^{\mathcal{B}}$). This implies that $(a, \mathrm{PD}^{\mathcal{B}}_\emptyset(\alpha')) \in \mathrm{LO}^{\mathcal{B}}_\emptyset(\alpha)$ but $(a, \mathrm{PD}^{\mathcal{B}}_\emptyset(\alpha')) \notin \mathrm{LO}^{\mathcal{B}}_\emptyset(\beta)$. Since $\mathrm{PD}^{\mathcal{B}}_\emptyset(\alpha) = \mathrm{PD}^{\mathcal{B}}_\emptyset(\beta) \neq \varepsilon$ (otherwise $\alpha = \varepsilon$ and we have no transition $\alpha \xrightarrow{a} \alpha'$), we get a contradiction with Prop. 22. □

6 Nondeterministic exponential-time algorithms

Bisimilarity problem (Theorem 2). As already mentioned, a proof of Theorem 2 is now obvious. A nondeterministic exponential-time algorithm, when

given a normed $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ and $A, B \in \mathcal{V}$, guesses an (at most exponential) representation of a pre-base $\mathcal{B}$, checks that it is a consistent base, and that $\mathrm{PD}^{\mathcal{B}}(A) = \mathrm{PD}^{\mathcal{B}}(B)$. Lemmas 20, 21, and 23 show that this nondeterministic algorithm indeed works in exponential time and that it has a successful run if, and only if, $A \sim B$ in $\mathcal{L}_{\mathcal{G}}$.

Regularity problem (Theorem 3). By using the results of previous sections we derive Lemma 25, which is the crux of a nondeterministic exponential-time algorithm deciding semantic finiteness of a given nBPA process. But we first recall a general notion of the bisimilarity quotient and note a simple fact.

Brbis-quotient. Given an LTS $\mathcal{L} = (\mathcal{S}, \mathcal{A}, (\stackrel{a}{\longrightarrow})_{a\in\mathcal{A}})$, the *quotient-LTS* (w.r.t. branching bisimilarity) is defined as $\mathcal{L}_{\sim} = (\{[s]; s \in \mathcal{S}\}, \mathcal{A}, (\stackrel{a}{\longrightarrow})_{a\in\mathcal{A}})$ where the states are the equivalence classes, hence $[s] = \{s' \mid s' \sim s\}$, and $[s_1] \stackrel{a}{\longrightarrow} [s_2]$ iff there are s_1', s_2' such that $s_1 \sim s_1'$, $s_2 \sim s_2'$, and $s_1 \stackrel{a}{\longrightarrow} s_2'$.

Since $\{(s, [s]) \mid s \in \mathcal{S}\} \cup \{([s], s) \mid s \in \mathcal{S}\}$ can be easily verified to be a branching bisimulation (on the disjoint union of $\mathcal{L}$ and $\mathcal{L}_{\sim}$), we have $s \sim [s]$.

Brbis-(in)finiteness. We say that a *state* r in an LTS is *brbis-finite* if r is branching bisimilar with some state in a finite LTS; otherwise r is *brbis-infinite*. We observe the next simple fact:

Proposition 24. *Given an LTS $\mathcal{L} = (\mathcal{S}, \mathcal{A}, (\stackrel{a}{\longrightarrow})_{a\in\mathcal{A}})$, a state $r \in \mathcal{S}$ is brbis-finite if, and only if, the set* $\textsc{reach}_{\sim}(r) = \{[s]; r \stackrel{w}{\longrightarrow} s \text{ for some } w \in \mathcal{A}^*\}$ *of equivalence classes reachable from r is finite.*

The next lemma is a consequence, when we recall that $\equiv^{\mathcal{B}} \subseteq \sim$ for any consistent base $\mathcal{B}$ (by Lemma 23) and $\sim = \equiv^{\mathcal{B}}$ for any intended base (as follows from Corollary 18(1)), which is also consistent (by Lemma 21).

Lemma 25. *Given an nBPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$, a configuration $\alpha \in \mathcal{V}^*$ is brbis-finite if, and only if, there exists a consistent base $\mathcal{B}$ such that the set*

$$\textsc{pd-reach}^{\mathcal{B}}(\alpha) = \{\mathrm{PD}^{\mathcal{B}}(\beta) \mid \alpha \stackrel{w}{\longrightarrow} \beta \text{ for some } w \in \mathcal{A}^*\}$$

of $(\mathcal{B}, \emptyset)$-prime-decompositions of configurations reachable from α is finite.

Proof. If $\mathcal{B}$ is a consistent base, then $\mathrm{PD}^{\mathcal{B}}(\beta_1) = \mathrm{PD}^{\mathcal{B}}(\beta_2)$ implies $\beta_1 \sim \beta_2$. Hence if $\textsc{pd-reach}^{\mathcal{B}}(\alpha)$ is finite, then we can reach only finitely many equivalence classes from α.

On the other hand, if α is branching bisimilar with a state in a finite LTS, then there are only finitely many equivalence classes reachable from α. For any intended (and consistent) base $\mathcal{B}$ we thus have that $\textsc{pd-reach}^{\mathcal{B}}(\alpha)$ is finite. □

Finishing a proof of Theorem 3. A nondeterministic algorithm deciding brbis-finiteness of a given α (for a given nBPA system $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$), can just guess an (at most exponential) representation of a pre-base $\mathcal{B}$, check that it is a consistent base, and that the set $\textsc{pd-reach}^{\mathcal{B}}(\alpha)$ is finite. Checking consistency

can be done in exponential time by Lemma 20; it thus remains to show how to test finiteness of $\text{PD-REACH}^{\mathcal{B}}(\alpha)$.

Given a (consistent) base $\mathcal{B}$, we say that (A, R) is a *PD-loop*, where $A \in \mathcal{V}$ and $R \in \text{DOM}(\mathcal{B})$, if there are $w \in \mathcal{A}^*$, $\beta \in \mathcal{V}^*$ such that $A \xrightarrow{w} A\beta$, $\text{PD}_R^{\mathcal{B}}(\beta) \neq \varepsilon$, and $\text{RED}^{\mathcal{B}}(\beta, R) = R$.

We now show that $\text{PD-REACH}^{\mathcal{B}}(\alpha)$ is infinite if, and only if, there exists $\delta = A\gamma$ such that δ is reachable from α, $R = \text{RED}^{\mathcal{B}}(\gamma)$ and (A, R) is a PD-loop; in this case we say that α *reaches a PD-loop*.

For the "if" direction note that then

$$\alpha \xrightarrow{v} A\gamma \xrightarrow{w} A\beta\gamma \xrightarrow{w} A\beta\beta\gamma \xrightarrow{w} A\beta\beta\beta\gamma \xrightarrow{w} \cdots$$

for some $v \in \mathcal{A}^*$. This path visits $\alpha_0 = A\gamma$, $\alpha_1 = A\beta\gamma$, $\cdots$, $\alpha_i = A\beta^i\gamma$, $\cdots$ where the set $\{\text{PD}_\emptyset^{\mathcal{B}}(\alpha_i) \mid i \in \mathbb{N}\}$ is infinite.

For the other direction, if $\text{PD-REACH}^{\mathcal{B}}(\alpha)$ is infinite, then there must be a path $\alpha = \alpha_0 \xrightarrow{a_1} \alpha_1 \xrightarrow{a_2} \alpha_2 \xrightarrow{a_3} \cdots$ where the set $\{\text{PD}_\emptyset^{\mathcal{B}}(\alpha_i) \mid i \in \mathbb{N}\}$ is infinite (which follows from König's Lemma). Then there obviously must be $i < j$ such that $\alpha_i = A\gamma$, $\alpha_j = A\beta\gamma$, and $A \xrightarrow{w} A\beta$ (for some $w \in \mathcal{A}^*$), where for $R = \text{RED}^{\mathcal{B}}(\gamma, \emptyset)$ we have $\text{RED}^{\mathcal{B}}(\beta\gamma, \emptyset) = \text{RED}^{\mathcal{B}}(\beta, R) = R$ and $\text{PD}_R^{\mathcal{B}}(\beta) \neq \varepsilon$; hence (A, R) is a PD-loop.

For finding PD-loops we can construct a directed graph as follows. We take the pairs (A, R) where $A \in \mathcal{V}$, $R \in \text{DOM}(\mathcal{B})$ as the vertices. Now we put an arc from (A, R) to (B, R') iff there is a rule $A \xrightarrow{a} \gamma_1 B \gamma_2$ of $\mathcal{G}$ such that $R' = \text{RED}^{\mathcal{B}}(\gamma_2, R)$; moreover, if there is such a case with $\text{PD}_R^{\mathcal{B}}(\gamma_2) \neq \varepsilon$, then we "colour" the arc as "blue". It is easy to check that (A, R) is a PD-loop iff there is a cycle from (A, R) to (A, R) in the graph that contains at least one blue arc.

Let us call a pair (B, R') *PD-infinite* if there is a path in the graph that starts in (B, R') and ends in a PD-loop (A, R). Now we observe that α reaches a PD-loop iff it can be written $\alpha = \beta_1 X \beta_2$ where (X, R) is PD-infinite for $R = \text{RED}^{\mathcal{B}}(\beta_2, \emptyset)$. An algorithm proving Theorem 3 is thus clear.

References

1. Baeten, J.C.M., Bergstra, J.A., Klop, J.W.: Decidability of bisimulation equivalence for processes generating context-free languages. J. ACM 40(3), 653–682 (1993)
2. Burkart, O., Caucal, D., Steffen, B.: An elementary bisimulation decision procedure for arbitrary context-free processes. In: Wiedermann, J., Hájek, P. (eds.) MFCS. Lecture Notes in Computer Science, vol. 969, pp. 423–433. Springer (1995)
3. Caucal, D., Huynh, D.T., Tian, L.: Deciding branching bimiliarity of normed context-free processes is in Σ_2^p. Inf. Comput. 118(2), 306–315 (1995)
4. Christensen, S., Hüttel, H., Stirling, C.: Bisimulation equivalence is decidable for all context-free processes. Inf. Comput. 121(2), 143–148 (1995)
5. Czerwinski, W.: Partially-commutative context-free graphs. Ph.D. thesis, University of Warsaw (2012)

6. Czerwinski, W., Hofman, P., Lasota, S.: Decidability of branching bisimulation on normed commutative context-free processes. Theory Comput. Syst. 55(1), 136–169 (2014)
7. Fu, Y.: Checking equality and regularity for normed BPA with silent moves. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M.Z., Peleg, D. (eds.) ICALP (2). Lecture Notes in Computer Science, vol. 7966, pp. 238–249. Springer (2013)
8. van Glabbeek, R.J., Weijland, W.P.: Branching time and abstraction in bisimulation semantics. J. ACM 43(3), 555–600 (1996)
9. Hirshfeld, Y., Jerrum, M., Moller, F.: A polynomial algorithm for deciding bisimilarity of normed context-free processes. Theor. Comput. Sci. 158(1&2), 143–159 (1996)
10. Hüttel, H.: Silence is golden: Branching bisimilarity is decidable for context-free processes. In: Proc. CAV'91. LNCS, vol. 575, pp. 2–12. Springer (1992)
11. Jančar, P.: Bisimilarity on basic process algebra is in 2-ExpTime (an explicit proof). Logical Methods in Computer Science 9(1) (2013)
12. Kiefer, S.: BPA bisimilarity is EXPTIME-hard. Inf. Process. Lett. 113(4), 101–106 (2013)
13. Mayr, R.: Weak bisimilarity and regularity of context-free processes is *exptime*-hard. Electr. Notes Theor. Comput. Sci. 96, 153–172 (2004)
14. Srba, J.: Roadmap of infinite results. In: Current Trends In Theoretical Computer Science, The Challenge of the New Century, vol. 2, pp. 337–350. World Scientific Publishing Co. (2004), updated version at http://users-cs.au.dk/srba/roadmap/
15. Yin, Q., Fu, Y., He, C., Huang, M., Tao, X.: Branching bisimilarity checking for PRS. In: Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.) ICALP (2). Lecture Notes in Computer Science, vol. 8573, pp. 363–374. Springer (2014)